



**HAL**  
open science

# Probabilistic Motion Modeling from Medical Image Sequences: Application to Cardiac Cine-MRI

Julian Krebs, Tommaso Mansi, Nicholas Ayache, Hervé Delingette

► **To cite this version:**

Julian Krebs, Tommaso Mansi, Nicholas Ayache, Hervé Delingette. Probabilistic Motion Modeling from Medical Image Sequences: Application to Cardiac Cine-MRI. STACOM 2019: Statistical Atlases and Computational Models of the Heart, Oct 2019, Shenzhen, China. hal-02239318v1

**HAL Id: hal-02239318**

**<https://inria.hal.science/hal-02239318v1>**

Submitted on 13 Aug 2019 (v1), last revised 23 Sep 2019 (v2)

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# Probabilistic Motion Modeling from Medical Image Sequences: Application to Cardiac Cine-MRI

Julian Krebs<sup>1,2</sup>, Tommaso Mansi<sup>1</sup>, Nicholas Ayache<sup>2</sup>, and Hervé Delingette<sup>2</sup>

<sup>1</sup> Siemens Healthineers, Digital Services, Digital Technology and Innovation, Princeton, NJ, 08540 USA

<sup>2</sup> Université Côte d’Azur, Inria, Epione Team, Sophia Antipolis, 06902 France

**Abstract.** We propose to learn a probabilistic motion model from a sequence of images. Besides spatio-temporal registration, our method offers to predict motion from a limited number of frames, useful for temporal super-resolution. The model is based on a probabilistic latent space and a novel temporal dropout training scheme. This enables simulation and interpolation of realistic motion patterns given only one or any subset of frames of a sequence. The encoded motion also allows to be transported from one subject to another without the need of inter-subject registration. An unsupervised generative deformation model is applied within a temporal convolutional network which leads to a diffeomorphic motion model – encoded as a low-dimensional motion matrix. Applied to cardiac cine-MRI sequences, we show improved registration accuracy and spatio-temporally smoother deformations compared to three state-of-the-art registration algorithms. Besides, we demonstrate the model’s applicability to motion transport by simulating a pathology in a healthy case. Furthermore, we show an improved motion reconstruction from incomplete sequences compared to linear and cubic interpolation.

## 1 Introduction

In medical imaging, an important task is to analyze temporal image sequences to understand physiological processes of the human body. Dynamic organs, such as the heart or lungs, are of particular interest to study as detected motion patterns are helpful for the diagnosis and treatment of diseases. Moreover, recovering the motion pattern allows to track anatomical structures, to compensate for motion, to do temporal super-resolution and motion simulation.

Motion is typically studied by computing pairwise deformations – the registration of each of the images in a sequence with a target image. The resulting dense deformation fields track moving structures from the beginning to the end of the sequence. Providing an invertible and smooth transformation, diffeomorphic registration algorithms such as the SyN algorithm [1] are especially suited for the registration of sequential images. One difficulty is to acquire temporally smooth deformations that are fundamental for tracking. That is why registration algorithms with a temporal regularizer have been proposed [7,8]. In the computer vision community, temporal video super-resolution and motion compensation are of related interest [4].

However, while these methods produce accurate dense deformations, they do not aim to extract intrinsic motion parameters crucial for building a comprehensive motion model useful for analysis tasks such as motion classification or simulation. Rohé et al. [9] proposed a parameterization, the Barycentric Subspaces, as a regularizer for cardiac motion tracking. Yang et al. [10] generated a motion prior using manifold learning from low-dimensional shapes.

We propose to learn a probabilistic motion model from image sequences directly. Instead of defining a parameterization explicitly or learning from pre-processed shapes, our model captures relevant motion features in a low-dimensional motion matrix in a generic but data-driven way. This learned latent space can be used to fill gaps of missing frames (motion reconstruction), to predict the next frames in the sequence or to generate an infinite number of new motion patterns given only one image (motion simulation). Motion can be also transported by applying the motion matrix on an image of another subject.

We apply a Bayesian inference method [5] in conjunction with a temporal regularizer [2] with explicit time dependence to learn the probabilistic motion encoding. To enforce temporal consistency, we introduce the self-supervised training scheme of temporal dropout sampling, which randomly leaves out input frames while trying to recover the complete motion sequence. The framework is learned in an unsupervised fashion from complete image sequences. Smooth, diffeomorphic and symmetric deformations are ensured by applying an exponentiation layer, spatio-temporal regularization and a symmetric local cross-correlation metric. Besides motion simulation, the model demonstrates state-of-the-art registration results for diffeomorphic tracking of cardiac cine-MRI. The main contributions are as follows:

- An unsupervised probabilistic motion model learned from image sequences
- A generative model using explicit time-dependent temporal convolutional networks trained with self-supervised temporal dropout sampling
- Demonstration of cardiac motion tracking, simulation, transport and temporal super-resolution

## 2 Methods

The motion observed in an image sequence with  $T + 1$  frames is typically described by deformation fields  $\phi_t$  between a moving image  $I_0$  and the fixed images  $I_t$  with  $t \in [1, T]$ . Inspired by the probabilistic deformation model of [6] based on conditional variational autoencoder (CVAE) [5], we define a motion model for temporal sequences. The model is conditioned on the moving image and parameterizes the set of diffeomorphisms  $\phi_t$  in a low-dimensional probabilistic space, the motion matrix  $z \in \mathbb{R}^{d \times T}$ , where  $d$  is the size of the deformation encoding per image pair. Each column's  $z_t$ -code corresponds to the deformation  $\phi_t$ . To take temporal dependencies into account,  $z_t$  is conditioned on all past and future time steps. To learn this temporal regularization directly from data, we apply Temporal Convolutional Networks [2] with explicit time dependence and temporal dropout sampling enforcing the network to fill time steps by looking at given past and future deformations. An illustration of the model is shown in Fig. 1a.

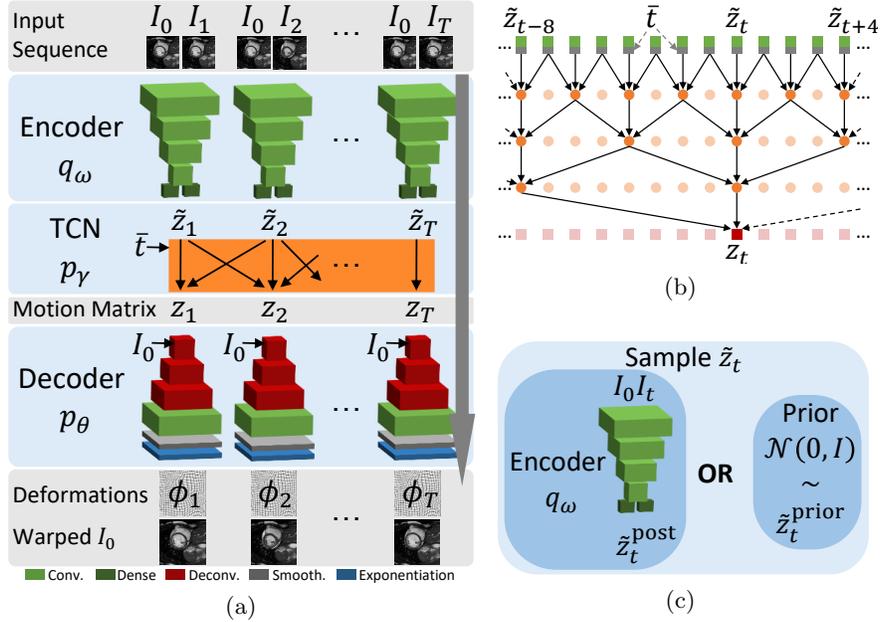


Fig. 1: Probabilistic motion model (a): The encoder  $q_\omega$  projects the image pair  $(I_0, I_t)$  to a low-dimensional deformation encoding  $\tilde{z}_t$  from which the temporal convolutional network  $p_\gamma$  (b) constructs the motion matrix  $z \in \mathbb{R}^{d \times T}$  conditioned on the normalized time  $\bar{t}$ . The decoder  $p_\theta$  maps the motion matrix to the deformations  $\phi_t$ . The temporal dropout sampling procedure (c) randomly chooses to sample  $\tilde{z}_t$  either from the encoder  $q_\omega$  or the prior distribution.

**Probabilistic Motion Model** Our motion model consists of three distributions. First, the encoder  $q_\omega(\tilde{z}_t|I_0, I_t)$  maps each of the image pairs  $(I_0, I_t)$  independently to a latent space denoted by  $\tilde{z}_t \in \mathbb{R}^d$ . Second, as the key component of temporal modeling, these latent vectors  $\tilde{z}_t$  are jointly mapped to the motion matrix  $z$  by conditioning them in all past and future time steps and on the normalized time  $\bar{t}$ :  $p_\gamma(z|\tilde{z}_{1:T}, \bar{t}_{1:T})$ . Finally, the decoder  $p_\theta(I_t|z_t, I_0)$  aims to reconstruct the fixed image  $I_t$  by warping the moving image  $I_0$  with the deformation  $\phi_t$ . This deformation  $\phi_t$  is extracted from the temporally regularized  $z_t$ -codes and conditioned on the moving image.

The distributions  $q_\omega, p_\gamma, p_\theta$  are approximated by three neural networks with trainable parameters  $\omega, \gamma, \theta$ . During training, a lower bound on the data likelihood is maximized with respect to a prior distribution  $p(\tilde{z}_t)$  of the latent space  $\tilde{z}_t$  (cf. CVAE [5]). The prior  $p(\tilde{z}_t)$  is assumed to follow a multivariate unit Gaussian distribution with spherical covariance  $I$ :  $p(\tilde{z}_t) \sim \mathcal{N}(0, I)$ . The objective function results in optimizing the expected log-likelihood  $p_\theta$  and the Kullback-Leibler (KL) divergence enforcing the posterior distribution  $q_\omega$  to be close to the prior  $p(\tilde{z}_t)$  for all time steps:

$$\sum_{t=1}^T \mathbb{E}_{z_t \sim p_\gamma(\cdot|\tilde{z}_{1:T}, \bar{t}_{1:T})} \left[ \log p_\theta(I_t|z_t, I_0) \right] - \text{KL} [q_\omega(\tilde{z}_t|I_0, I_t) \| p(\tilde{z})]. \quad (1)$$

Unlike the traditional CVAE model, the temporal regularized  $z_t$ -code is used in the log-likelihood term  $p_\theta$  instead of the  $\tilde{z}_t$ . We model  $p_\theta$  as a symmetric local cross-correlation Boltzmann distribution with the weighting factor  $\lambda$ . Encoder and decoder weights are independent of the time  $t$ . Their network architecture consists of convolutional and deconvolutional layers with fully-connected layers for mean and variance predictions in the encoder part [5]. We use an exponentiation layer for the stationary velocity field parameterization of diffeomorphisms [6], a linear warping layer and diffusion-like regularization with smoothing parameters  $\sigma_G$  in spatial and  $\sigma_T$  in temporal dimension.

### Temporal Convolutional Networks with Explicit Time Dependence

Since the parameters of encoder  $q_\omega$  and decoder  $p_\theta$  are independent of time, the temporal conditioning  $p_\gamma$  plays an important role in merging information across different time steps. In our work, this regularization is learned by Temporal Convolutional Networks (TCN). Consisting of multiple 1-D convolutional layers with increasing dilation, TCN can handle input sequences of different lengths. TCN have several advantages compared to recurrent neural networks such as a flexible receptive field and more stable gradient computations [2].

The input of the TCN is the sequence of  $\tilde{z}$  concatenated with the normalized time  $\bar{t} = t/T$ . Providing the normalized time explicitly, provides the network with information on where each  $\tilde{z}$  is located in the sequence. This supports the learning of a motion model from data representing the same type of motion with varying sequence lengths. The output of the TCN is the regularized motion matrix  $z$ . We use non-causal instead of causal convolutional layers to also take future time steps into account. We follow the standard implementation using zero-padding and skip connections. Each layer contains  $d$  filters. A schematic representation of our TCN is shown in Fig. 1b.

**Training with Temporal Dropout Sampling** Using Eq. 1 for training could lead to learning the identity transform  $z \approx \tilde{z}$  in the TCN  $p_\gamma$  such that deformations of the current time step are independent of past and future time steps. To avoid this and enforce the model to search for temporal dependencies during the training, we introduce the concept of temporal dropout sampling (TDS). In TDS, some of the  $\tilde{z}_t$  are sampled from the prior distribution  $p(\tilde{z})$  instead of only sampling from the posterior distribution  $q_\omega(\tilde{z}_t|I_0, I_t)$  as typical for CVAE. At the time steps the prior has been used for sampling, the model has no knowledge of the target image  $I_t$  and is forced to use the temporal connections within the TCN in order to minimize the objective.

More precisely, at each time step  $t$ , a sample from the prior distribution  $\tilde{z}_t^{\text{prior}} \sim p(\tilde{z}_t)$  is selected instead of a posterior sample  $\tilde{z}_t^{\text{post}} \sim q_\omega(\tilde{z}_t|I_0, I_t)$  using a binary Bernoulli random variable  $r_t$ . All independent Bernoulli random variables  $r \in \mathbb{R}^T$  have the success probability  $\delta$ . The latent vector  $\tilde{z}_t$  can be defined as:

$$\tilde{z}_t = r_t * \tilde{z}_t^{\text{prior}} + (1 - r_t) * \tilde{z}_t^{\text{post}}. \quad (2)$$

Fig. 1c illustrates the TDS procedure. At test time, for each time step independently, one can either draw  $\tilde{z}_t$  from the prior or take the encoder’s prediction.

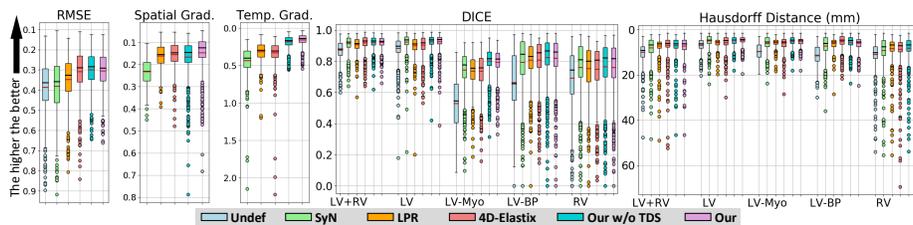


Fig. 2: Tracking results showing RMSE, spatial and temporal gradient of the displacement fields, DICE scores and Hausdorff distances. The proposed algorithms (Our and Our w/o TDS) show slightly higher registration accuracy and temporally smoother deformations than the state-of-the-art algorithms SyN [1], LPR [6] and 4D-Elastix [7].

### 3 Experiments

We evaluate our motion model on 2-D cardiac MRI-cine data. First, we demonstrate accurate temporal registration by evaluating motion tracking and compensation of the cardiac sequence, taking the end-diastolic (ED) frame as the moving image  $I_0$ . Stabilization, is accomplished by warping all frames  $I_t$  to the ED frame. Pair-wise registration results are presented for ED-ES (end-systolic) frame pairs. Second, we present motion transport, motion sampling and reconstruction with a limited number of frames.

**Data** We used 334 short-axis sequences acquired from different hospitals including 150 sequences from the Automatic Cardiac Diagnosis Challenge 2017 (ACDC [3]). The remaining cases were obtained from the EU FP7-funded project MD-Paedegree, mixing congenital heart diseases with healthy and pathological images from adults. The sequence length  $T$  varied from 13 to 35 frames. We used the 100 cases from ACDC that contain ED-ES segmentation information for testing while the remaining sequences were used for training. All slices were resampled with a spacing of  $1.5 \times 1.5$  mm and cropped to a size of  $128 \times 128$  pixels.

**Implementation Details** The encoder consisted of 4 convolutional layers with strides (2, 2, 2, 1) and the decoder had 3 deconvolutional and 1 convolutional layer before the exponentiation and warping layers (Fig. 1a). The TCN consists of four layers with dilation (1, 2, 4, 8) and a kernel size of 3 (cf. Fig. 1b). The regularization parameters  $\sigma_G$  and  $\sigma_T$  were set to 3 mm respectively 1.5. The loss weighting factor  $\lambda$  was chosen empirically as  $6 \cdot 10^4$ . The deformation encoding size  $d$  was set to 32. The dropout sampling probability  $\delta$  was 0.5. We applied the Adam optimization method with a learning rate of 0.00015 and a batch size of one. We performed data augmentation on-the-fly by randomly shifting, rotating, scaling and mirroring images. We implemented the model in Tensorflow with Keras. The training time was 15h on a NVIDIA GTX TITAN X GPU.

#### 3.1 Registration: Tracking and Motion Compensation

We compare our model with and without the temporal dropout sampling (Our w/o TDS) with three state-of-the-art methods: SyN [1], the learning-based al-

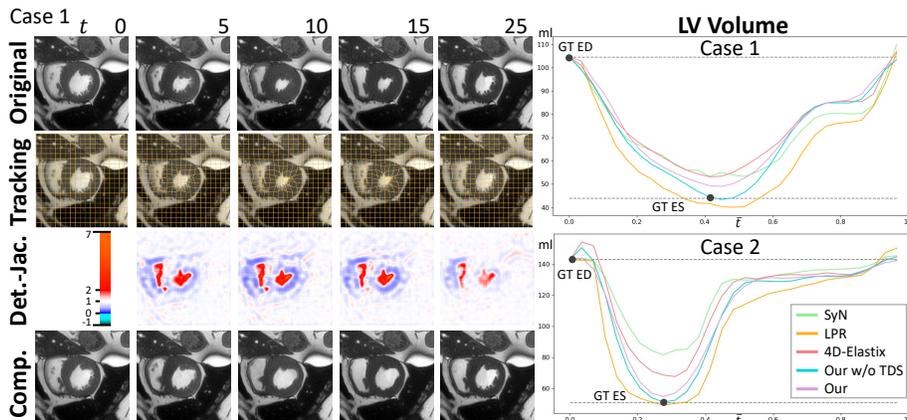


Fig. 3: Left: Visual results of one test sequence, showing tracking, the Jacobian determinant (Det.-Jac.) and motion compensation (Comp.). Right: LV volumes in ml for this sequence and another one (ground truth ED/ES volumes marked with points).

algorithm (LPR, 2-D single-scale version [6]) and the b-spline-based 4D algorithm in elastix [7]. In Fig. 2, tracking results are visualized for the test data taking sequences of all slices where segmentation in ED and ES frames are available, resulting in 677 sequences. We report the root mean square error (RMSE), the spatial (Spatial Grad.) and temporal gradient (Temp. Grad) of the displacement fields for evaluating the smoothness of the resulting deformations. Our model shows spatially and temporally smoother deformations. We also report DICE scores and 95%-tile Hausdorff distances (HD in mm) on five anatomical structures: myocardium (LV-Myo) and epicardium (LV) of the left ventricle, left bloodpool (LV-BP), right ventricle (RV) and LV+RV. Note, that DICE scores and HD were evaluated on ED-ES frame pairs only, while the other metrics were computed on the full sequence. The proposed method showed improved mean DICE scores and smaller mean HD of 84.6%, 6.2mm (w/o TDS: 84.7%, 6.1mm) compared to SyN, LPR, 4D-Elastix with (82.7%, 7.0mm), (82.1%, 6.6mm) respectively (83.7%, 6.3mm%). Compared to training without temporal dropout sampling, HD and DICE scores show minimal differences. However, only TDS offers consistent motion simulation and temporal interpolation. Also, deformation regularity is improved using TDS.

Visual results for one case are shown on the left of Fig. 3. On the figure’s right side the LV volume curves computed by warping the ED mask are plotted. One can see that the SyN algorithm underestimates big deformations. Both, the volume and the gradient metrics show smoother deformations for both versions of our motion model compared to the SyN, LPR and 4D-Elastix algorithms.

### 3.2 Sampling, Sequence Reconstruction and Motion Transport

We then evaluate motion sampling, reconstruction and transport. We extract simulated reconstructed motion patterns if we provide our model with different subsets of images from the original sequence. In the time steps without frames,

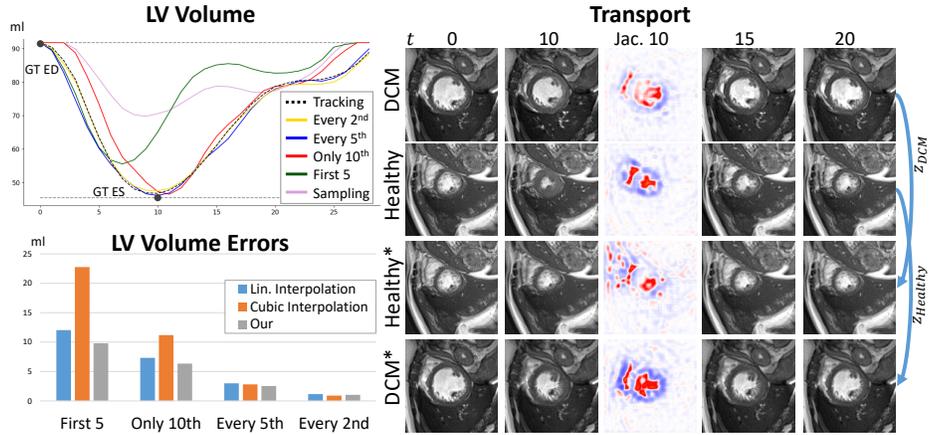


Fig. 4: Left Top: LV volume curves from simulated and reconstructed motion by providing only a subset of frames and predicting the complete motion sequence from them. We provided only the first frame  $I_0$  (sampling), every second, every fifth, the first five frames or only the tenths frame of the same sequence. Left Bottom: Mean volume errors with respect to the tracking volumes of all 677 testing cases comparing our sampling procedure with linear interpolation between velocity fields of the given frames. Right: The motion matrix  $z$  from a pathological (first row: dilated myopathy DCM) and a healthy subject (second row) are transported from one to the other (bottom rows).

the motion matrix  $z$  is created by randomly sampling  $\tilde{z}_t$  from the prior distribution as depicted in Fig. 1c (take same  $\tilde{z}_t$  for all slices of one volume). Here, we choose sampling over interpolation in  $\tilde{z}$ -space to remain an uncertain, a probabilistic, estimation of the interpolated deformations. The left side of Fig. 4 shows LV volume curves and reconstruction errors if every second, every fifth, only the 10th, no frame (sampling) or the first 5 frames are provided besides the moving frame (ED). The LV volume errors are computed on all test sequences by taking the mean absolute differences between sampled and tracking volumes. We compare with linear and cubic interpolation of velocities (extracted from tracking) between given times. One can see that our model’s relative performance in recovering the LV volume increases when fewer frames are provided. Given the first five frames or only one additional frame (10th frame), the model estimates the motion consistently with plausible cardiac motion patterns for the missing time steps. In the cases of providing every second and every fifth frame, our method performs equally good or marginally better. Note, the motion simulation given only the ED frame (sampling) does not overlap with the original motion, which is not intended. Nevertheless, one can see cardiac specific motion patterns such as the plateau phase before the atrial systole.

Motion can be transported by taking the motion matrix  $z$  from one sequence and apply it on the ED frame  $I_0$  of another sequence. The right side of Fig. 4 shows the transport of a pathological motion to a healthy subject and vice versa. The resulting simulated motion shows similar heart contractions and motion

characteristics as the originating motion while the transported deformations are adapted to the target image without requiring explicit inter-subject registration.

## 4 Conclusions

In this paper, we presented an unsupervised approach for learning a motion model from image sequences. Underlying motion factors are encoded in a low-dimensional probabilistic space, the motion matrix, in which each column represents the deformation between two frames of the sequence. Our model demonstrated accurate motion tracking and motion reconstruction from missing frames, which can be useful for shorter acquisition times and temporal super-resolution. We also showed motion transport and simulation by using only one frame. For future work, we aim to further explore the spatial coherence between slices for 3-D applications and the influence of using different training datasets (pathological and non-pathological) on the learned motion matrix.

**Disclaimer:** This feature is based on research, and is not commercially available. Due to regulatory reasons its future availability cannot be guaranteed.

## References

1. Avants, B.B., Epstein, C.L., Grossman, M., Gee, J.C.: Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain. *Medical image analysis* **12**(1), 26–41 (2008)
2. Bai, S., Kolter, J.Z., et al.: An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. *arXiv preprint arXiv:1803.01271* (2018)
3. Bernard, O., et al.: Deep learning techniques for automatic MRI cardiac multi-structures segmentation and diagnosis: Is the problem solved? *IEEE Transactions on Medical Imaging* **37**(11), 2514–2525 (2018)
4. Caballero, J., Ledig, C., Aitken, A., et al.: Real-time video super-resolution with spatio-temporal networks and motion compensation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 4778–4787 (2017)
5. Kingma, D.P., et al.: Semi-supervised learning with deep generative models. In: *Advances in Neural Information Processing Systems*. pp. 3581–3589 (2014)
6. Krebs, J., Delingette, H., Mailhé, B., et al.: Learning a probabilistic model for diffeomorphic registration. *IEEE Transactions on Medical Imaging* **38** (Feb 2019)
7. Metz, C., Klein, S., Schaap, M., van Walsum, T., Niessen, W.J.: Nonrigid registration of dynamic medical imaging data using nd+ t b-splines and a groupwise optimization approach. *Medical image analysis* **15**(2), 238–249 (2011)
8. Qin, C., Bai, W., et al.: Joint learning of motion estimation and segmentation for cardiac mr image sequences. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 472–480. Springer (2018)
9. Rohé, M.M., Sermesant, M., Pennec, X.: Low-dimensional representation of cardiac motion using barycentric subspaces: A new group-wise paradigm for estimation, analysis, and reconstruction. *Medical image analysis* **45**, 1–12 (2018)
10. Yang, L., Georgescu, B., Zheng, Y., et al.: Prediction based collaborative trackers (pct): A robust and accurate approach toward 3d medical object tracking. *IEEE transactions on medical imaging* **30**(11), 1921–1932 (2011)